The effect of subtracting a constant from all payoffs in a hawk–dove game: experimental evidence of loss aversion in strategic behavior

Nick Feltovich*
University of Aberdeen Business School
Edward Wright Building, Dunbar Street
Aberdeen AB24 3QY, UK
n.feltovich@abdn.ac.uk

January 19, 2010

Abstract

Economists and psychologists have documented patterns of individual decision–making behavior (e.g., loss aversion) whereby losses and gains are treated differently. However, there has been little evidence of such patterns in multi–player games. We report results showing the strongest evidence we know of that this phenomenon is present in games. Experimental subjects play two hawk–dove games that are identical up to a constant; in one, all payoffs are positive, while in the other, payoffs are negative iff both players choose "hawk". Under both fixed–pairs and random matching, differences between the games are substantial, significant, and consistent with loss aversion.

Journal of Economic Literature classifications: D81, C72, C73.
Keywords: experiment, game theory, behavioral economics, chicken, loss aversion, loss avoidance.

*This research was partly funded by the University of Aberdeen Business School. I thank Miguel Costa-Gomes, Erika Seki, Joe Swierzbinski, and Qin Xiao for helpful suggestions and comments.

1 Background

Game–theoretic solution concepts imply that a change in payoff levels—modifying a game by adding a constant to all payoffs—should not affect behavior as long as payoffs reflect players' preferences (i.e., they are equivalent to expected utilities). However, when the entries in a "payoff matrix" are actually monetary gains or losses, as in economics experiments with human subjects, the robustness of behavior to such changes becomes an empirical question. Research has found that payoff levels can indeed have an effect on behavior, particularly when the *signs* of payoffs change as a result: gains turn to losses or vice versa. Kahneman and Tversky (1979) document survey evidence suggesting that losses weigh more heavily than equal–sized gains, and name this phenomenon "loss aversion". Many examples of loss aversion in laboratory experiments and in the field have since been found (Camerer 2004). However, nearly all have involved individual decision–making tasks; by contrast, evidence of loss aversion in strategic decision making has been hard to find (see Section 2 for a review).

The objective of this paper is to add to the discussion of whether, and how, strategic behavior is affected by differential treatment of gains and losses, by reporting the strongest evidence we have seen thus far of such an effect in a strategic situation—specifically, results of a human–subjects experiment involving two hawk–dove games (see Figure 1) played under either a fixed–pairs or a random–matching protocol. We are agnostic about the source of the effect we find—whether it comes from loss aversion or

Figure 1: The hawk–dove games used in the experiment

		Player 2				Player 2	
		Dove	Hawk			Dove	Hawk
Player	Dove	160,160	80,200		Dove	120,120	40,160
1	Hawk	200,80	20,20		Hawk	160,40	–20,–20
		High payoffs (HIGH)				Low payoffs (LOW)	

some other phenomenon (such as loss avoidance, described in Section 2)—so to save space, we will often abuse terminology somewhat by using "loss aversion" as a shorthand term for any pattern of behavior in which losses weigh more heavily than gains, irrespective of the underlying cause.[1]

2 Previous literature

Here, we present a brief review of previous research into the effects of changing payoff levels (for a more thorough review, see Feltovich, Iwasaki, and Oda (2008)). Ido Erev and colleagues have conducted several studies examining payoff–level effects in repeated individual–decision problems, typically finding that subjects are quicker to learn an optimal choice under limited payoff information when both gains and losses are

[1]We stress that we are not arguing that all such patterns of behavior are equivalent to loss aversion (e.g., Feltovich, Iwasaki, and Oda (2008) show that neither loss aversion nor loss avoidance implies the other, although they sometimes make the same prediction), but rather that we did not design our experiment to distinguish among these various alternatives.

possible than when only gains or only losses are possible.[2] Erev, Bereby–Meyer, and Roth (1999) examined behavior in two versions of a repeated 2x2 constant–sum stage game with a unique Nash equilibrium in mixed strategies. Payoffs in the games differed only by subtraction of a constant, but one game had losses possible while the other did not. Erev, Bereby–Meyer, and Roth found that subject choices were more consistent with fictitious–play learning (that is, more likely to optimize against the historical distribution of opponent choices) when losses were possible than when they were not. On the other hand, Rapoport and Boebel (1992) failed to find a payoff–level effect in two versions of a repeated 5x5 constant–sum game with a unique mixed–strategy equilibrium. Unlike Erev, Bereby–Meyer, and Roth, though, Rapoport and Boebel do not affect the signs of any payoffs with their manipulation of payoff levels, suggesting that not all payoff–level changes matter for behavior—only those that change the signs of payoffs.

Another way of examining payoff–level effects is in market experiments, in which subjects play the role of firms facing sunk costs that vary across treatments. Results in these experiments have depended sensitively on characteristics of the market institution that was used: when the institution is one that tends strongly toward equilibrium, sunk costs have typically had no effect (e.g., Kachelmeier (1996) for a double auction and Waller, Shapiro, and Sevcik (1999), for a multi–firm posted–price market), while for institutions with weaker tendencies toward equilibrium, sunk costs can have an effect (e.g., Offerman and Potters (2003) and Buchheit and Feltovich (2010), both of which used price–setting duopolies with capacity constraints).

Several researchers have used coordination games with Pareto–ranked Nash equilibria to study the effects of changing the payoff level. Cachon and Camerer (1996) implemented payoff–level changes in a multi–player median–effort game by varying a (mandatory) sunk cost, and found that subjects behaved differently when some actions led to guaranteed losses compared with a baseline case where all could yield gains; they named this tendency "loss avoidance". However, Cachon and Camerer found no evidence of loss avoidance in a third treatment where subjects were told their own sunk cost but not those of the other players, so they concluded that the effect arose only indirectly, from subjects choosing actions under the belief that their opponents exhibited loss avoidance—that is, due to the strategic complementarities in the game.[3] Rydval and Ortmann (2005) examined two pairs of stag–hunt games, with one member of each pair differing from the other by a constant. They found behavior consistent with loss avoidance in one pair, but not in the other.[4] Feltovich, Iwasaki, and Oda (2010) draw a distinction between the loss avoidance studied by these previous researchers (which they term "certain–loss avoidance") and "possible–loss avoidance", which is a tendency to avoid actions leading to possible losses when alternatives yielding certain gains are available. Feltovich, Iwasaki, and Oda found evidence of both types of loss avoidance in repeated stag–hunt games, but the effect was small when subjects had complete information about payoffs. Finally, Johnson, Myagkov, and Orbell (2005) found that decisions to enter a prisoners' dilemma game were

[2] See, for example, Barkan, Zohar, and Erev (1998), Bereby–Meyer and Erev (1998), and Erev, Bereby–Meyer, and Roth (1999). There was also earlier work that didn't look specifically at payoff–level effects, but gave results suggesting that changes in payoff levels might affect behavior (e.g., Siegel and Goldstein (1959) and Siegel, Siegel, and Andrews (1964)). Finally, we note here that while Kahneman and Tversky (1979)'s prospect theory includes loss aversion—thus allowing for payoff–level changes to affect behavior—they do not specifically test the effects of such changes.

[3] Indeed, Camerer (2003) concludes that "loss–avoidance is a selection principle applied to infer what others will do...not a principle of individual choice" (p. 393).

[4] Rydval and Ortmann suggest that the higher scale of payoffs in the latter pair led to the effects of loss aversion counteracting those of loss avoidance.

sensitive to whether payoffs were framed as gains or losses, but that decisions within the game—conditional on choosing to play—were unaffected.

3 Methods

Our examination of loss aversion involves the two games shown in Figure 1. These games are identical up to a constant; the low–payoff game is obtained from the high–payoff game by subtracting the same number (40) from all payoffs. Nash equilibria in both games are (Hawk, Dove), (Dove, Hawk), and a mixed–strategy equilibrium in which each player chooses Dove with probability $p = \frac{3}{5}$.

Hawk–dove games have a noteworthy advantage over other games that have been used for examining loss aversion, most of which had the property of *strategic complementarities*: a strategy becomes *more* attractive (relative to alternatives) as the likelihood of other players choosing it increases. One drawback of using games with strategic complementarities is that they make it difficult to disentangle choices of a strategy due to preference for that strategy (e.g., loss aversion) from choices due to a belief that one's opponent prefers that strategy (e.g., belief that others are loss averse). A second issue arises from the fact that loss aversion can only have an effect on behavior when genuine strategic uncertainty is present, that is, when players believe that multiple opponent actions have non–negligible likelihood. When strategic complementarities are present, pure–strategy Nash equilibria will be symmetric, so in the neighborhood of an equilibrium, there will not be genuine strategic uncertainty in the game (since nearly everyone will be choosing the same pure strategy). As a result, it becomes difficult to detect any effect from loss aversion, especially after subjects have gained experience playing the game; at best, the researcher might hope to find a second–order effect in differential likelihoods of the various pure–strategy equilibria.[5]

In the hawk–dove game, by contrast, the two strategies are *strategic substitutes*: a strategy becomes *less* attractive as the likelihood of opponents choosing it increases. This property, in conjunction with other features of our experimental design, allows us to avoid both of the shortcomings mentioned above. First, it becomes straightforward to distinguish between direct effects via preferences and indirect effects via beliefs; if (for example) a player believes her opponent due to loss aversion is relatively likely to choose Dove, she should be *less*, not more, likely to choose Dove herself, whereas her own loss aversion would make her *more* likely to choose Dove. Second, one of our treatments uses a one–population random–matching design (in a given round, a subject is equally likely to be matched with any other subject). Since the three Nash equilibria involve population aggregate choice probabilities of Dove of $\frac{1}{2}$, $\frac{1}{2}$, and $\frac{3}{5}$ (for the two pure and one mixed equilibria respectively)—quite far from unanimous choice of a single strategy in each case—it should nearly always be true that there is genuine strategic uncertainty in that treatment, and thus that loss aversion should have visible effects.[6]

[5] We are not claiming that loss aversion cannot affect riskless choice in general; indeed, there is some evidence that it does (see, for example, Tversky and Kahneman (1991) and Gächter, Johnson, and Herrmann (2007)). We are simply noting that in the particular riskless decisions that can arise in these particular games (for example, the choice of safe versus risky strategy in a stag hunt, when the opponent is believed to be choosing the safe strategy with certainty), loss aversion will have no effect on behavior.

[6] The one–population random–matching protocol, combined with symmetry of the games and our lack of subject labels (e.g., all subjects saw themselves as Player 1 in the experiment), makes it unlikely that pairs could successfully coordinate on either of the asymmetric pure–strategy Nash equilibria, and thus that the population could converge to either of these equilibria. One could thus argue that the relevant solution concept is evolutionarily stable strategy (Maynard Smith (1974)), not Nash equilibrium. The only evolutionarily stable strategy of either game is the mixed–strategy Nash equilibrium.

As an example of how loss aversion might affect behavior in these games, suppose a player's payoff is equivalent to her monetary payment in all cases except in the low–payoff game's (Hawk, Hawk) outcome, whose monetary payment of –20 is instead treated as a payoff of -20λ, where $\lambda \geq 0$. In that case, λ can be regarded as a measure of that individual's attitude toward losses; $\lambda = 1$ means losses and gains are treated the same, while larger values represent loss aversion. Then, the mixed–strategy equilibrium probability of choosing Dove is $p(\lambda) = \frac{2+\lambda}{4+\lambda}$, which is strictly increasing in λ. For $\lambda = 1$, we have $p = \frac{3}{5}$, so that behavior is the same as in the high–payoff game (where loss aversion does not come into play), whereas for $\lambda > 1$, there will be a higher equilibrium probability of Dove in the low–payoff game than in the high–payoff game. This logic motivates our hypothesis:

Hypothesis 1 *Choices of Dove should be* more *likely in the low–payoff version of the hawk–dove game than in the high–payoff version.*

Subjects in the experiment played each hawk–dove game twenty times; that is, the game was varied within–subject. We used two matching mechanisms—random matching and fixed pairs—which we varied between–subjects. Under random matching, subjects were matched randomly in each round, with any other subject equally likely to be the opponent in a given round. Under fixed pairs, subjects were matched randomly in the first round of each game, but played versus the same subject for all rounds of a game.

The use of fixed pairs leads to more complex theoretical predictions, compared to random matching, in two ways. First, under fixed pairs, the one–population argument made in Note 6 does not apply, so that the pure–strategy stage–game equilibria become more likely outcomes. Second, folk theorems for finitely–repeated games allow play of the (Dove, Dove) action profile in all but the last two rounds of a game in subgame perfect equilibrium.[7] However, these points are equally valid in both high– and low–payoff versions of the game.

Subjects were given complete payoff information in the experiment. In order to disguise the similarity between the two hawk–dove games, subjects played three additional games, all of which were versions of a stag–hunt game. (See Figure 2.) The ordering of games always had the form SH–HD–SH–HD–SH, but we

Figure 2: The stag–hunt games used in the experiment

High payoffs (SHH)

		Player 2	
		Risky	Safe
Player 1	Risky	360,360	40,260
	Safe	260,40	260,260

Medium payoffs (SHM)

	Player 2	
	Risky	Safe
Risky	220,220	–100,120
Safe	120,–100	120,120

Low payoffs (SHL)

	Player 2	
	Risky	Safe
Risky	80,80	–240,–20
Safe	–20,–240	–20,–20

varied the ordering of the three stag–hunt games and the two hawk–dove games across sessions. (See the text surrounding Table 2 for a discussion of order effects.)

The experimental sessions were conducted at the Scottish Experimental Economics Laboratory (SEEL) at the University of Aberdeen. Sessions typically comprised 10–20 subjects. Some larger sessions were split

[7] Other subgame perfect equilibria exist with (Dove, Dove) profiles sustained for fewer rounds. For a discussion of folk theorem results for finitely–repeated games, see Fudenberg and Tirole (1992), pp. 165–168.

into sub–sessions—so that subjects in one sub–session did not interact with those in others—in order to increase the number of independent observations. Subjects were primarily undergraduate students from University of Aberdeen, and were recruited from a database of people expressing interest in participating in experiments. No one took part in more than one session.

The experiment was run on networked computer terminals, using the z–Tree experiment software package (Fischbacher 2007).[8] At the beginning of a session, subjects were seated in a single room and given a set of written instructions (samples are provided in the appendix). The instructions were also read aloud to the subjects, in an attempt to make the rules common knowledge, after which the first round of play began. The instructions stated that twenty rounds of each game would be played; when the twentieth round was completed, it was publicly announced that the game was going to change, and that subjects would be re–matched (even in the fixed–pairs treatment). Subjects were asked not to communicate with other subjects, so the only interactions took place via the computer program. Subjects were given no identifying information about the opposing player, in an attempt to minimize incentives for reputation building and other supergame effects in the random–matching treatment. Also, rather than using potentially biasing terms like "opponent" or "partner" for the other player, we used the neutral though somewhat cumbersome "player matched to you" and similar phrases.

Each round began with a screen announcing that a new round is beginning, and when appropriate, that the game and/or opponent has changed. At that time, subjects were prompted to make their choices for that round; to avoid demand effects, the actions were displayed in neutral terms as "X" and "Y". After all subjects had made their choices, each was given the following feedback: own choice, opponent choice, and own payoff. Subjects were not told the opponent's payoff, though they did receive enough information to easily determine this if they wished. After all subjects clicked a button on the screen to continue, the session proceeded to the next round.

At the end of a session, subjects were paid in cash, privately and individually. Each subject's payment included a £8 show–up fee; additionally, one round from each game was randomly chosen, and the subject was paid his/her earnings in those rounds, at an exchange rate of £1 per 100 points.[9] Total earnings for subjects participating in a session averaged about £14, with most earning £9–£18, for a session that lasted 45–75 minutes.

4 Experimental results

Thirteen sessions were conducted: 6 sessions (7 sub–sessions) with a total of 90 subjects for the random–matching treatment, and 7 sessions (8 sub–sessions) with 86 subjects for the fixed–pairs treatment. Table 1 shows some aggregate results: frequencies of Dove choices in round 1, in rounds 16–20 (the last five rounds), and over all rounds of a game. Also shown are significance results from nonparametric statistical tests of

[8]The experiment programs, as well as the raw data, are available from the author upon request.

[9]An implication of our payment scheme is that subjects cannot lose money overall, as was required for human–subjects approval (though it was possible to earn less than the show–up fee). As a result, there might be some question as to whether negative entries in payoffs should be expected to be viewed by subjects as actual losses, as opposed to reduced gains. We simply note that, to the extent that subjects did not perceive these as "real" losses, our experiment is a conservative test for loss aversion.

Table 1: Aggregate Dove choice frequency, by game (HIGH or LOW) and matching treatment

	Random matching		Fixed pairs	
Rounds	HIGH	LOW	HIGH	LOW
1	.676	.822***	.721	.837**
16–20	.533	.605**	.700	.737*
All	.562	.650***	.724	.798*

** (**, ***): Higher in LOW than HIGH at the 10% (5%, 1%) level. See text for details.*

differences between the high– and low–payoff game.[10] The table shows clearly higher frequencies of Dove choices in the low–payoff game than in the high–payoff game, consistent with our hypothesis. The difference between the games is apparent in the first round, over the last five rounds, and overall, and is statistically significant in all cases (at the 5% level or better in the random–matching treatment and at the 10% level or better in the fixed–pairs treatment).[11]

More information about aggregate behavior is given by Figure 3, which shows the time paths of Dove choice frequencies in the experiment. While the strength of strategic considerations in this environment are

Figure 3: Round–by–round frequency of Dove choice

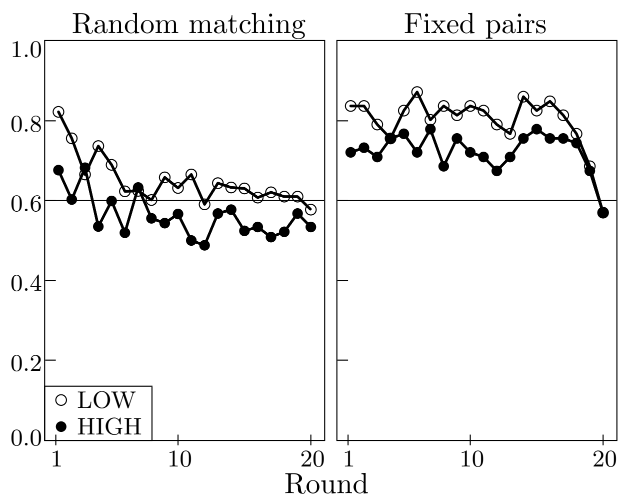

evident, the figure also shows systematic differences between the two hawk–dove games. In the random–matching treatment, Dove choice frequencies are initially higher than the equilibrium levels: substantially so for the low–payoff game, less so for the high–payoff game. These frequencies trend downward in both games, approaching the mixed–strategy equilibrium level in the low–payoff game, while actually staying below that level throughout the last ten rounds of the high–payoff game.

Under fixed pairs, on the other hand, Dove choices in both games occur substantially more often than in equilibrium—with no visible time trend—until nearly the end, when Dove choice frequencies sharply

[10]When testing first–round choices, we use the chi–square test for paired replicates with individual–level data. For the last five rounds and the 20–round averages, we use a one–tailed Wilcoxon signed–ranks test for paired replicates with sub–session–level data. See Siegel and Castellan (1988) for descriptions of these tests.

[11]The difference over all rounds of the fixed–pairs treatment just misses being significant at the 5% level ($p \approx 0.0547$).

decline, reaching roughly equilibrium levels in the last round. These trajectories seem to reflect partially successful attempts by players to cooperate, along the lines of folk theorems for finitely–repeated games (as noted in Section 3).

We next look at some probit regression results, allowing us to ascertain the treatment effect while controlling for other factors that may affect behavior, such as the order in which the games were played. We estimate two models, one for the fixed–pairs treatment and one for the random–matching treatment. The dependent variable in each regression is an indicator for a Dove choice in the current round. In order to determine the effect of the payoff level (and thus test for loss aversion), we use as our primary explanatory variable an indicator for the LOW game; in addition, we include the products of this indicator with the round number and its square to allow for the possibility of a time–varying effect. As controls for possible order effects, we include an indicator for the LOW–HIGH ordering of the two hawk–dove games, as well as indicators for which particular stag–hunt game preceded the current hawk–dove game (in case the previous game played has any effect on behavior in the current game). Finally, we included indicators for the round number and its square. The regressions were performed on the sub–sample of hawk–dove plays (we left out the stag–hunt data) using Stata (version 10), and incorporated individual–subject random effects.

Table 2: Results of probit regressions with random effects (std. errors in parentheses)

Dependent variable: Dove choice (round t)	random–matching treatment ($N = 3600$)	fixed–pairs treatment ($N = 3440$)
constant	0.482***	0.686
	(0.157)	(0.420)
LOW	0.419**	0.290
	(0.165)	(0.185)
LOW · Round	–0.021	0.047
	(0.034)	(0.040)
LOW · Round^2	–0.0006	–0.003*
	(0.0016)	(0.002)
p–value (joint significance of three LOW variables)	< 0.001***	< 0.001***
Round	–0.060**	0.050*
	(0.23)	(0.027)
Round^2	–0.002*	–0.002**
	(0.001)	(0.001)
LOW–HIGH game ordering	–0.005	–0.099
	(0.077)	(0.330)
prev. game was SHM	0.088	0.238
	(0.068)	(0.328)
prev. game was SHL	0.047	–0.065
	(0.082)	(0.640)
–ln(L)	2073.862	1438.546

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

Table 2 shows coefficients and standard errors for each variable, as well as log likelihoods for each regression. The results suggest that subject behavior is nonstationary; indeed, the coefficients for the four variables containing the round number or its square are jointly significant in both treatments ($p < 0.001$). On the other hand, only equivocal evidence exists for order effects: none of the order–effects variables are individually significant, though in the fixed–pairs treatment, these three variables are jointly significant at the 1% level.

Regarding loss aversion, we note first that one of its implications—joint significance of the three LOW variables (LOW, LOW·Round, and LOW·Round^2)—is seen in the table for both regressions. Next, in order to measure how loss aversion affects choice frequencies, we estimate the incremental effect of switching the game from HIGH to LOW. In round t, the total effect of the LOW variables on the argument of the normal c.d.f. used in the probit model is given by $\beta_{\text{LOW}} + \beta_{\text{LOW}\cdot\text{Round}} \cdot t + \beta_{\text{LOW}\cdot\text{Round}^2} \cdot t^2$ (where β_Y is the coefficient of the variable Y). So, the incremental effect of the LOW game rather than the HIGH game in round t has the form

$$\Phi\left(\bar{X} \cdot B + \beta_{\text{LOW}} + \beta_{\text{LOW}\cdot\text{Round}} \cdot t + \beta_{\text{LOW}\cdot\text{Round}^2} \cdot t^2\right) - \Phi\left(\bar{X} \cdot B\right), \tag{1}$$

where Φ is the normal c.d.f., $\bar{X}$ is the row vector of the other right–hand–side variables' values (which we set to an appropriate sample or subsample mean), and B is the column vector of their coefficients. If subjects exhibit loss aversion in round t, the sign of this expression should be positive for that value of t.

Graphs of this expression, for both treatments, are displayed in Figure 4, with point estimates for each round shown as circles and corresponding 90% confidence intervals shown as line segments.[12] The

Figure 4: Estimates of the effect of LOW game on Dove choice probability (vs. HIGH game)

(Circles represent point estimates; line segments represent 90% confidence intervals)

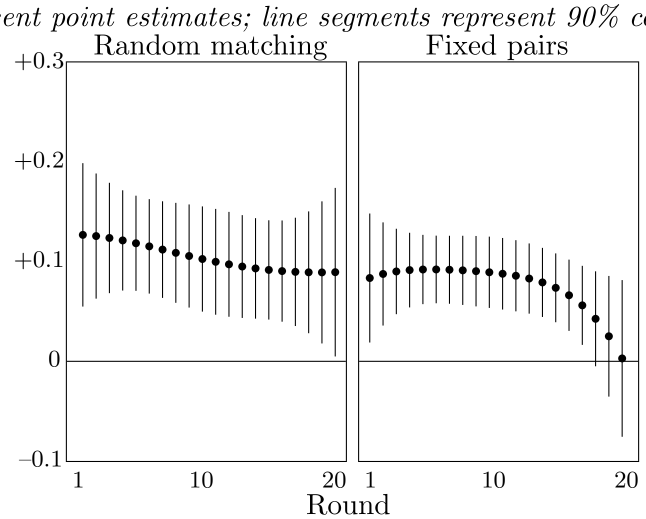

figure confirms what we had seen in the aggregate results. The point estimate for the incremental effect of the LOW game is always positive, and the corresponding 90% confidence interval lies entirely above zero (so that a one–sided hypothesis test would yield significance at the 5% level) for all rounds of the random–matching treatment, and all but the last three rounds of the fixed–pairs treatment. The positive effect of the LOW game in both treatments is in the direction predicted by loss aversion.

5 Summary

We have conducted an experiment involving two versions of a hawk–dove stage game, in order to examine the possibility of differential treatment of gains and losses in strategic decision making. The games we use are strategically identical to each other, differing only by a constant. The constant was chosen in such a

[12]Note that we use 90% confidence intervals rather than the usual 95% confidence intervals here. Since loss aversion makes a directional prediction, our rejection regions are one–tailed. Use of 90% confidence intervals gives us 5% rejection regions on the appropriate side.

way that in our “high–payoff” game, all payoffs are gains, while in our “low–payoff” game, the (Hawk, Hawk) outcome yields a loss, with all other outcomes still yielding gains. Loss aversion (or Feltovich, Iwasaki, and Oda’s (2008) possible–loss avoidance) predicts a higher probability of Dove choices in the low–payoff game relative to the high–payoff game. This prediction is borne out in the experimental results, where a statistically significant effect is observed under both random–matching and fixed–pairs matching protocols. Under random matching, the overall frequency of Dove choices is slightly lower in the high–payoff game than in the mixed–strategy Nash equilibrium, but this frequency rises by about 15% (about 9 percentage points) in the low–payoff game. Under fixed pairs, frequencies of Dove choices are higher in both games than under random matching (probably reflecting attempts at cooperation), but again, higher in the low–payoff game than in the high–payoff game, by about 10% (about 7 percentage points). The differences between high– and low–payoff games are visible—for both treatments—in the first round (before subjects have had any experience in the game), in late rounds (after subjects have learned about the strategic environment and the behavior of others in the game), and in session aggregates.

It is worth emphasizing the striking nature of this result. Previous research into the differential treatment of gains and losses has found strong evidence of an effect in individual decision making, but little such evidence in strategic behavior. If such a pattern of results were consistently seen, it would be possible to argue that loss aversion and its ilk are rather fragile phenomena, and easily weeded out when market discipline is present—even if the discipline is as weak as that in a two–player game. We see, however, that this is not the case. Not only have we observed behavior consistent with loss aversion, but we have seen it persisting in an environment where strategic considerations appear to be strong (as evidenced by the time paths seen in Figure 3). Indeed, the evidence we find for loss aversion in this experiment likely understates the true effect; as already mentioned, while an individual’s own loss aversion will make Dove choices more likely in the low–payoff game than in the high–payoff game, a belief that her opponent avoids losses will push her behavior in the opposite direction (because Hawk and Dove choices are strategic substitutes), leading to a overall effect smaller than would have come from the individual’s loss aversion by itself. In conclusion, we urge researchers to consider loss aversion when formulating predictions in strategic environments where losses are possible.

References

Barkan, R., D. Zohar, and I. Erev (1998), “Accidents and decision making under uncertainty: a comparison of four models,” *Organizational Behavior and Human Decision Processes* 74, pp. 118–144.

Bereby–Meyer, Y. and I. Erev (1998), “On learning to become a successful loser: a comparison of alternative abstractions of learning processes in the loss domain,” *Journal of Mathematical Psychology* 42, pp. 266–286.

Buchheit, S. and N. Feltovich (2010), “Experimental evidence of a sunk–cost paradox: a study of pricing behavior in Bertrand–Edgeworth duopoly,” forthcoming, *International Economic Review.*

Cachon, G.P. and C.F. Camerer (1996), “Loss avoidance and forward induction in experimental coordination games,” *Quarterly Journal of Economics* 111, pp. 165–194.

Camerer, C.F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton, NJ.

Camerer, C.F. (2004), "Prospect theory in the wild: evidence from the field," in *Advances in Behavioral Economics*, C.F. Camerer, G. Loewenstein, and M. Rabin, eds., Princeton University Press, Princeton, NJ, pp. 148–161.

Erev, I., Y. Bereby–Meyer, and A.E. Roth (1999), "The effect of adding a constant to all payoffs: experimental investigation and implications for reinforcement learning models," *Journal of Economic Behavior and Organization* 39, pp. 111–128.

Feltovich, N., A. Iwasaki, and S.H. Oda (2008), "Payoff levels, loss avoidance, and equilibrium selection in the Stag Hunt: an experimental study," working paper, University of Aberdeen.

Fischbacher, U. (2007), "z–Tree: Zurich toolbox for ready–made economic experiments," *Experimental Economics* 10, pp. 171–178.

Fudenberg, D. and J. Tirole (1992), *Game Theory*, MIT Press, Cambridge, MA.

Gächter, S., E.J. Johnson, and A. Herrmann (2007), "Individual–level loss aversion in riskless and risky choices," CeDEx working paper 2007–02, University of Nottingham.

Johnson, T., M. Myagkov, and J. Orbell (2005), "A bias toward loss aversion in the choice to enter risky cooperative games," working paper, http://www.allacademic.com/meta/p_mla_apa_research_citation/0/4/0/0/5/p40057_index.html.

Kahneman, D. and A. Tversky (1979), "Prospect theory: an analysis of decision making under uncertainty," *Econometrica* 47, pp. 273–297.

Maynard Smith, J. (1974), "The theory of games and the evolution of animal conflicts," *Journal of Theoretical Biology* 47, pp. 209–221.

Rapoport, A. and R.B. Boebel (1992), "Mixed strategies in strictly competitive games: a further test of the minmax hypothesis," *Games and Economic Behavior* 4, pp. 261–283.

Rydval, O. and A. Ortmann (2005), "Loss avoidance as selection principle: evidence from simple stag–hunt games," *Economics Letters* 88, pp. 101–107.

Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.

Siegel, S. and D.A. Goldstein (1959), "Decision–making behavior in a two choice uncertain outcome situation," *Journal of Experimental Psychology* 57, pp. 37–42.

Siegel, S., A.E. Siegel, and J.M. Andrews (1964), *Choice, Strategy, and Utility*, McGraw–Hill, New York.

Tversky, A. and D. Kahneman (1991), "Loss aversion in riskless choice: a reference–dependent model," *Quarterly Journal of Economics* 106, pp. 1039–1061.

Tversky, A. and D. Kahneman (1992), "Advances in prospect theory: cumulative representation of uncertainty," *Journal of Risk and Uncertainty* 5, pp. 297–323.

Appendix: experiment instructions

Note: Text in square brackets: shown only in fixed-pairs treatment. Text in parentheses: shown only in random-matching treatment.

Instructions

You are about to participate in an experiment in the economics of decision-making. Please read these instructions carefully, as the amount of money you earn will depend on how well you understand them. All money you earn will be paid to you privately in cash at the end of the experimental session. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment.

This experimental session consists of five different games. Each game will be played for 20 rounds. Each game will be played by you and another player. The player matched to you will be assigned randomly at the beginning of each [game, but will remain the same for all rounds of a game] (round). You will not be told the identity of the player matched to you, nor will he/she be told your identity—even after the end of the session.

The payoff tables: The number of points you earn in each round will depend on the choice you make and the choice made by the player matched to you. The payoff table for each game you play will be shown on your computer screen. In each round, both you and the player matched to you will have a choice between two possible actions, which will be called X and Y. Your action, together with the action chosen by the player matched to you, determines one of the four boxes in the payoff table. That box will show both your payoff and the payoff of the player matched to you.

Example: This is an example of a payoff table. You will not use this actual payoff table in the experiment, but the tables you will use will have a similar structure.

		Other player action	
		X	Y
Your action	X	Your payoff: 25 Other player payoff: 50	Your payoff: 125 Other player payoff: 150
	Y	Your payoff: 150 Other player payoff: 125	Your payoff: 50 Other player payoff: 25

Sequence of play in a round: The sequence of play in a round is as follows.

(1) [If it is the first round of a game, the] (The) computer randomly matches you to another participant.

(2) You and the player matched to you play the game. You choose an action, either X or Y. The player matched to you also chooses an action, either X or Y. Both of you make your choice without knowing the other's choice.

(3) The round ends. You receive the following information: your own choice, the choice made by the player matched to you, and your own payoff.

Payments: At the end of the experimental session, one of the rounds is chosen randomly from each game you have played, for a total of five rounds. You will be paid the total number of points you earned in those five rounds, at an exchange rate of £0.01 per point. Each participant additionally receives £8 for completing the session. Payments are made in cash at the end of the session.

Are there any questions?